A Strong Separation for Adversarially Robust ℓ_0 Estimation for Linear Sketches

Elena Gribelyuk[1], Honghao Lin[2], David P. Woodruff[2], Huacheng Yu[1], Samson Zhou[3]

[1] Princeton University, [2] Carnegie Mellon University, [3] Texas A&M University

Standard Streaming Model

- **Input:** Elements of a stream π, which arrive sequentially one at a time (*worst-case*, fixed in advance).
- **Output:** At the end of the stream, A outputs an approximation of a given function of π.
- **Goal:** A should use space *sublinear* in the size m of the input stream π.

Adversarially Robust Streaming

- **Input:** Elements of a stream π, which arrive sequentially and *adversarially.*
- **Output:** At each time t, A receives an update u_t, updates its internal state, and returns a *current* estimate r_t, which is recorded by the *adversary*.

"Future updates may depend on previous updates"

- **Question:** can we still design algorithms that use sublinear space?

Distinct Elements Estimation

- Given a stream π of m elements from $[n]$, let f_i denote frequency of element i.
- Let F_0 be the number of distinct elements: $F_0 = |\{i : f_i \neq 0\}|$
- **Goal**: Given a stream π of m elements from $[n]$ and an accuracy parameter ε, output a $(1+\varepsilon)$-approximation to F_0
- $\Theta(\frac{1}{\varepsilon^2} + \log n)$ space in standard streaming model.
  - Must use randomization to achieve sublinear space!
  - If the stream updates are adaptive, the adversary may (over time) learn something about the internal randomness.

Main Result

Theorem 1: There is a constant $\varepsilon = \Omega(1)$ so that any linear sketch giving a $(1+\varepsilon)$-approximation to F_0 on an adversarial insertion-deletion stream that uses $r < n^c$ rows, for a constant $c > 0$, can be broken in $\tilde{O}(r^8)$ queries.

Overview of our Approach

Construct adaptive attack for the gap ℓ_0 norm problem, defined below.

Definition (Gap ℓ_0 Norm Promise Problem):
Given input $x \in \mathbb{Z}^n$, decide whether $|x|_0 \geq \beta n$ or $|x|_0 \leq \alpha n$, for constants $0 < \alpha < \beta < 1$. If neither holds, return 0/1 arbitrarily.

High-level intuition:

- For query x, A will observe Ax.
- Some coordinates of the input vector are *significant*, i.e., learned well by the sketching matrix A, but most of them are not…
  - Ex. If sketching matrix A has a row e_i, A will observe $\langle e_i, x\rangle = x_i$ exactly.

Definition (significant coordinate):
Coordinate i is *significant* if there exists $y \in \mathbb{R}^r$ such that

$$(\mathrm{FRAC}(y^\top A)_i)^2 \geq \frac{1}{s}\sum \left(\mathrm{FRAC}(y^\top A)_j\right)^2$$

- WLOG, pre-process A to obtain a new matrix A', which separates the significant coordinates (sparse part) and insignificant coordinates (dense part).

$$A' = \begin{bmatrix} S \\ D \end{bmatrix}$$

Attack Outline:

1. Iteratively identify the significant coordinates and set them to zero in all future queries. Our attack algorithm is inspired by the *interactive fingerprinting code problem* (defined below).
2. After we have learned all significant coordinates, the query algorithm must rely on the other coordinates, for which the sketch Ax only has "small" information.
3. Finally, we design a hard distribution family $\mathcal{D}$ over $[-R, \ldots, R]$ for the dense part, such that
   - For $D_p \in \mathcal{D}$ with $p \in [\alpha, \beta]$, we have $\Pr_{X \sim D_p}[X = 0] = p$
   - For any $q, p \in [a, b]$, the total variation distance between Dx_p and Dx_q is small, i.e., $\frac{1}{\mathrm{poly}(n)}$.

Interactive Fingerprinting Code

- An algorithm $\mathcal{P}$ selects a secret set $S \subset [N], |S| = n$ of coordinates unknown to the fingerprinting code $\mathcal{F}$
- $\mathcal{F}$ must identify S by making adaptive queries $c^t \in \{0,1\}^N$
- For each query c^t, $\mathcal{P}$ must distinguish between the case that $c^t = 0^n$ versus $c^t = 1^n$.
- BUT: $\mathcal{P}$ can only observe c_i^t for $i \in S$.
- There exists an interactive fingerprinting code with length $\tilde{O}(n^2)$ [SteinkeUllman15]

Conclusions and Future Work

- We also provide efficient $poly(r)$ length adaptive attacks against linear sketches over $\mathbb{F}_p$ and $\mathbb{R}$.
- Open question: lower bounds against general sketches?